Tunneling Internet protocols inside QUIC

draft-piraux-intarea-quic-tunnel
draft-piraux-intarea-quic-tunnel-session
draft-piraux-intarea-quic-tunnel-tcp

Maxime Piraux, Olivier Bonaventure, Adi Masputra

This work was partially supported by the MQUIC project

Content

- **Introduction to QUIC**
- The tunnel mode

  draft-piraux-intarea-quic-tunnel
- The tunnel session mode

  draft-piraux-intarea-quic-tunnel-session
- The stream mode

  draft-piraux-intarea-quic-tunnel-tcp
- Conclusion

Introduction to QUIC

- Provides services equivalent to TCP+TLS atop UDP
- 1-RTT Authenticated Handshake
  - 0-RTT “session resumption”
- All application data and most control data is encrypted
  - Immune to middleboxes interference, and likely to pass through given its adoption for the web
- Two manners of conveying application data
  - Streams: reliable, in-order, uni- and bi-directional bytestreams
  - Datagrams: unreliable messages
- QUIC seems like a good fit for an alternative to IPSec, (D)TLS tunnels

Content

- Introduction to QUIC
- **The tunnel mode**

  draft-piraux-intarea-quic-tunnel
- The tunnel session mode

  draft-piraux-intarea-quic-tunnel-session
- The stream mode

  draft-piraux-intarea-quic-tunnel-tcp
- Conclusion

Reference environment

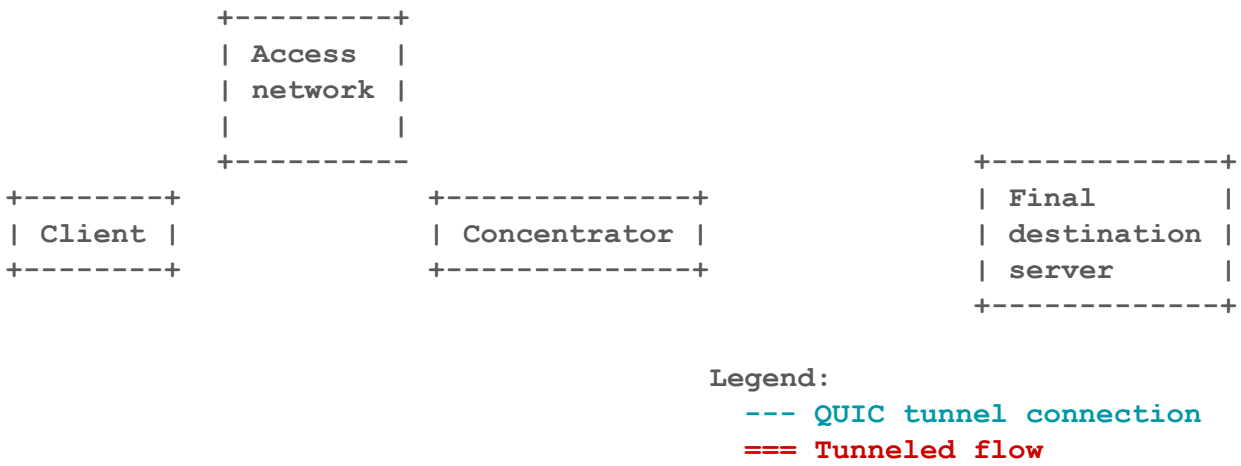

- Client uses a Concentrator to convey its traffic over the access network.

Reference environment

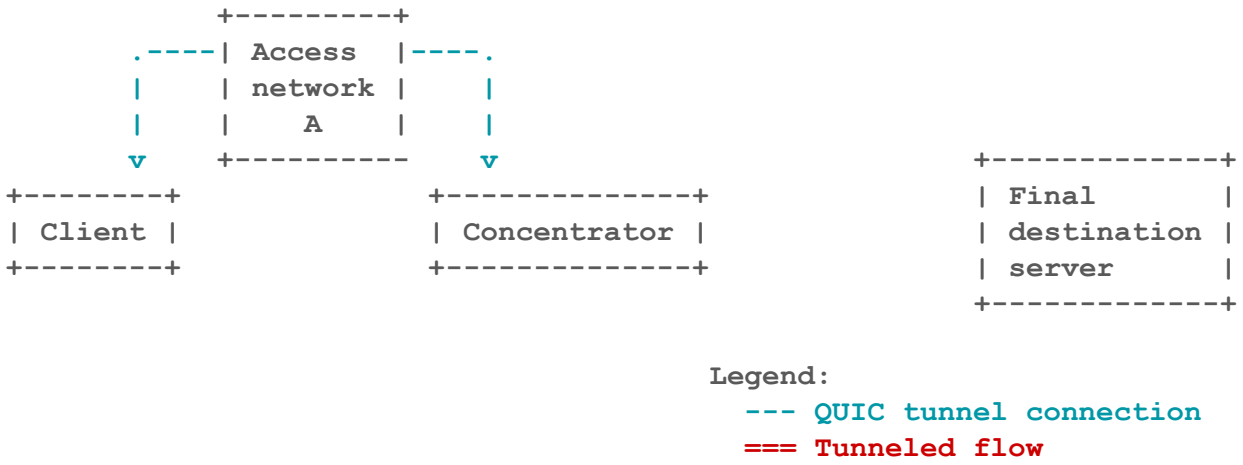

- Client uses a Concentrator to convey its traffic over the access network.
- Client sends all its packets to the Concentrator over a QUIC connection.

Reference environment

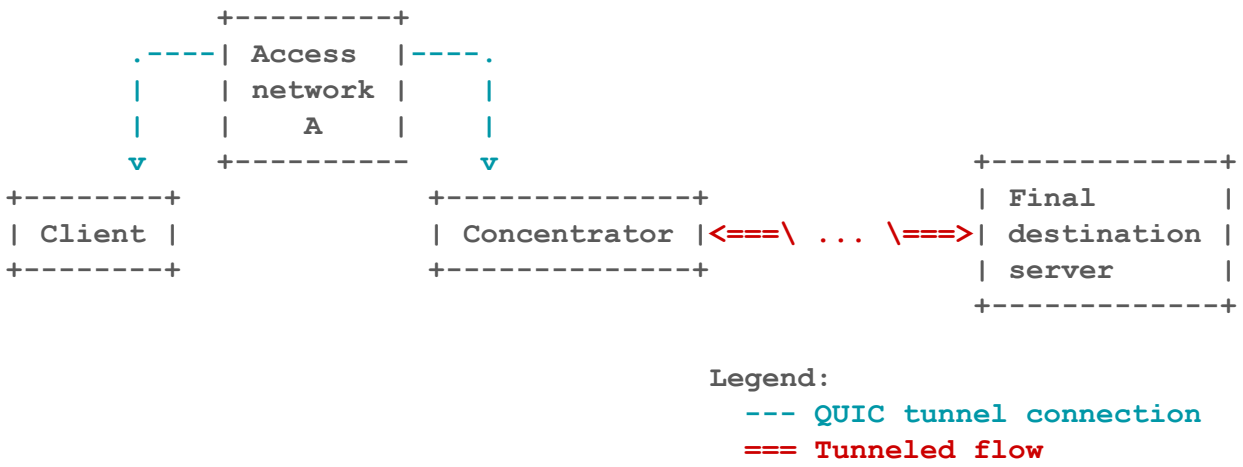

- Client uses a Concentrator to convey its traffic over the access network.
- Client sends all its packets to the Concentrator over a QUIC connection.
- The Concentrator forwards them to their final destination.
- Returning traffic destined to the Client is sent over the QUIC connection

The tunnel mode

- Negotiated using the “qt” ALPN token.
- Packets are transmitted inside QUIC DATAGRAM frames.
- Out-of-band signalling is used to negotiate the type of exchanged packets.

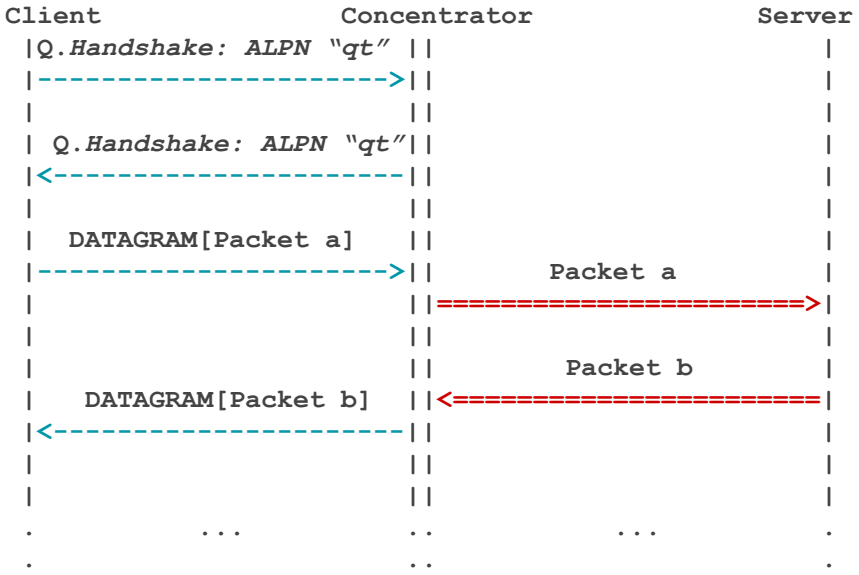

The tunnel mode

- `Access Report TLV (AR TLV)` allows to report access network availability status.
- The Client can signal when the network is unstable to stop incoming data.
- Later, it can resume the use of the QUIC tunnel connection in the same manner.

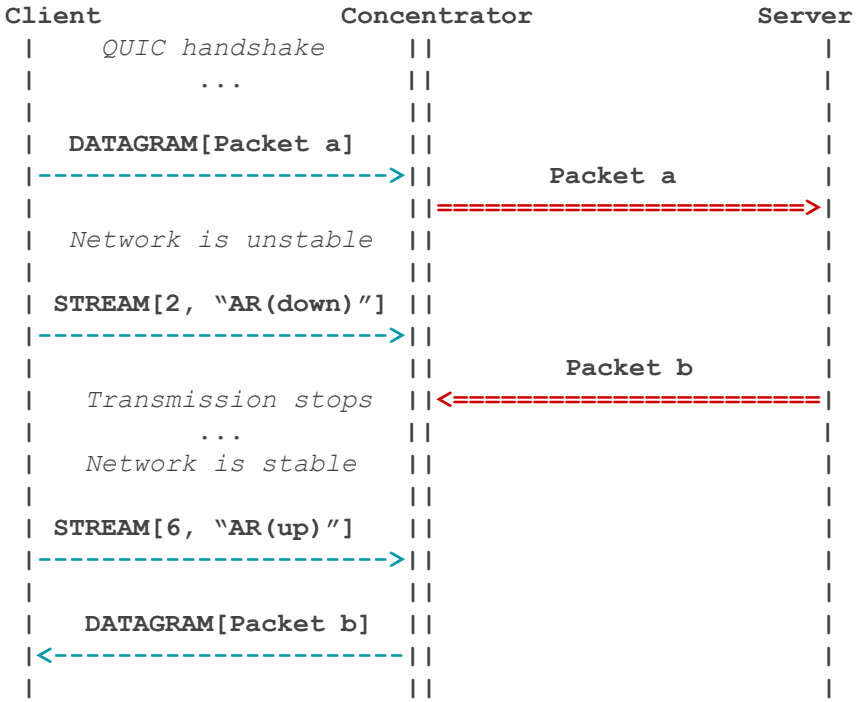

Content

- The tunnel mode

  draft-piraux-intarea-quic-tunnel
- **The tunnel session mode**

  draft-piraux-intarea-quic-tunnel-session
- The stream mode

  draft-piraux-intarea-quic-tunnel-tcp
- Conclusion

Reference environment

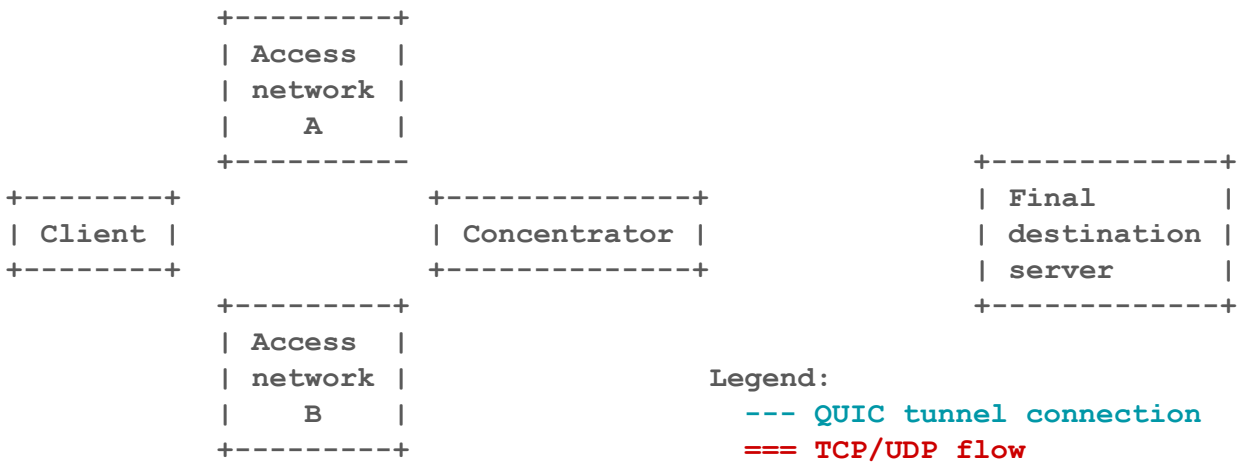

- Client uses a Concentrator to convey its traffic over the access network.
- Client sends all its packets to the Concentrator over a QUIC connection.
- The Concentrator forwards them to their final destination.
- Client is often multihomed and/or multistack, e.g. WiFi and 5G, IPv4 and IPv6.
- Client would like to leverage both access networks.
  - e.g. for load-sharing or fail-over

Reference environment

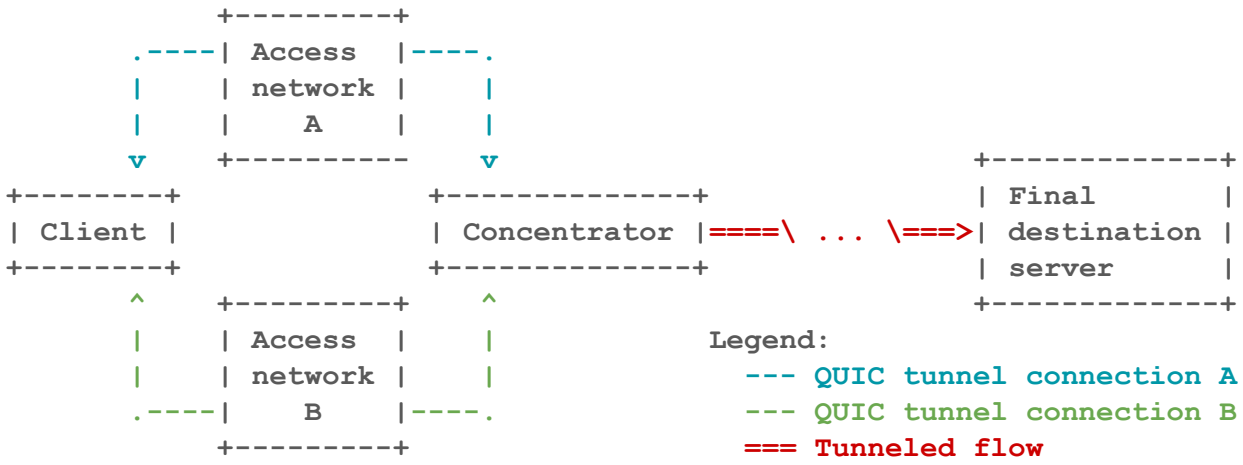

- Client uses a Concentrator to convey its traffic over the access network.
- Client sends all its packets to the Concentrator over a QUIC connection.
- The Concentrator forwards them to their final destination.
- Client is often multihomed and/or multistack, e.g. WiFi and 5G, IPv4 and IPv6.
- Client would like to leverage both access networks.
  - e.g. for load-sharing or fail-over

Reference environment

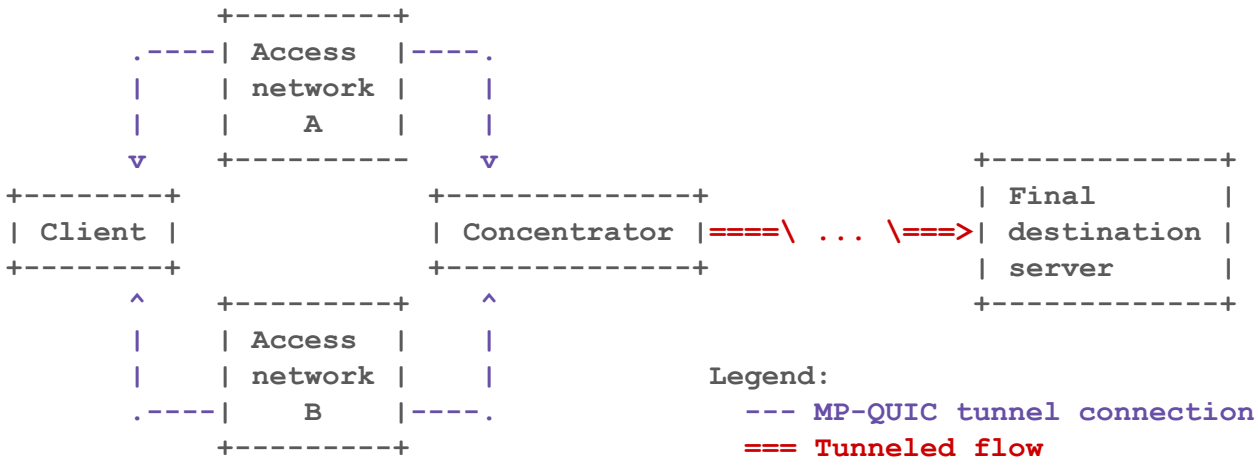

- Client uses a Concentrator to convey its traffic over the access network.
- Client sends all its packets to the Concentrator over a QUIC connection.
- The Concentrator forwards them to their final destination.
- Client is often multihomed and/or multistack, e.g. WiFi and 5G, IPv4 and IPv6.
- Client would like to leverage both access networks.
  - e.g. for load-sharing or fail-over

The tunnel session mode

- Connections can be grouped into a QUIC tunnel session.
- Allows coordinating packet reordering across connections.
- An optional opaque value indicates the QoS requested for each connection.

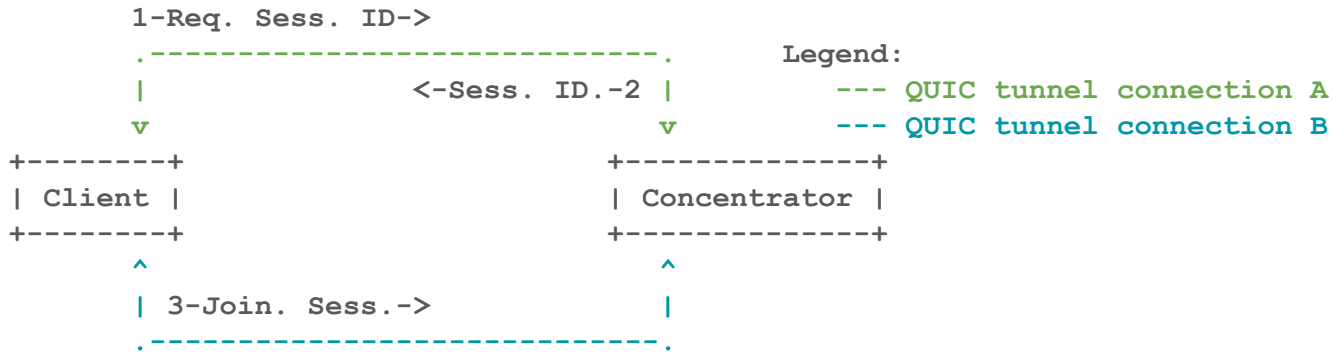

The tunnel session mode

- Explicitly identifies L2 or L3 packets exchanged over the QUIC connection.
- Packets are encoded inside a QUIC DATAGRAM frame using the format:

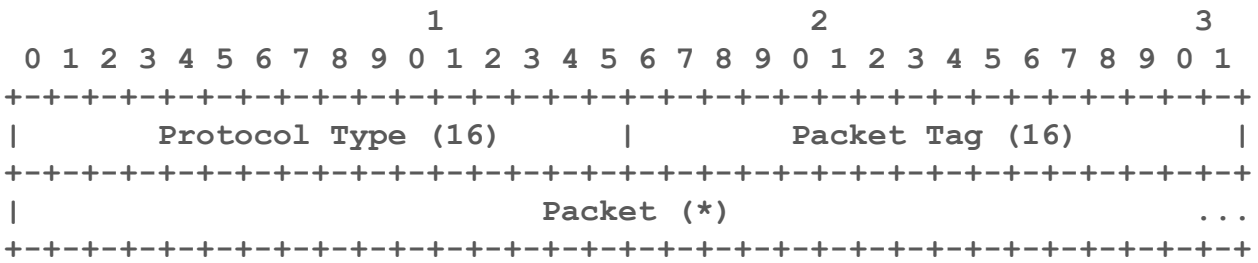

- Protocol Type: The protocol type, as in the “ETHER TYPES" IANA registry.
- Packet Tag: An opaque value. It can be used for reordering.

The tunnel session mode

- Connections can be grouped into a QUIC tunnel session.
- Allows coordinating packet reordering across connections.

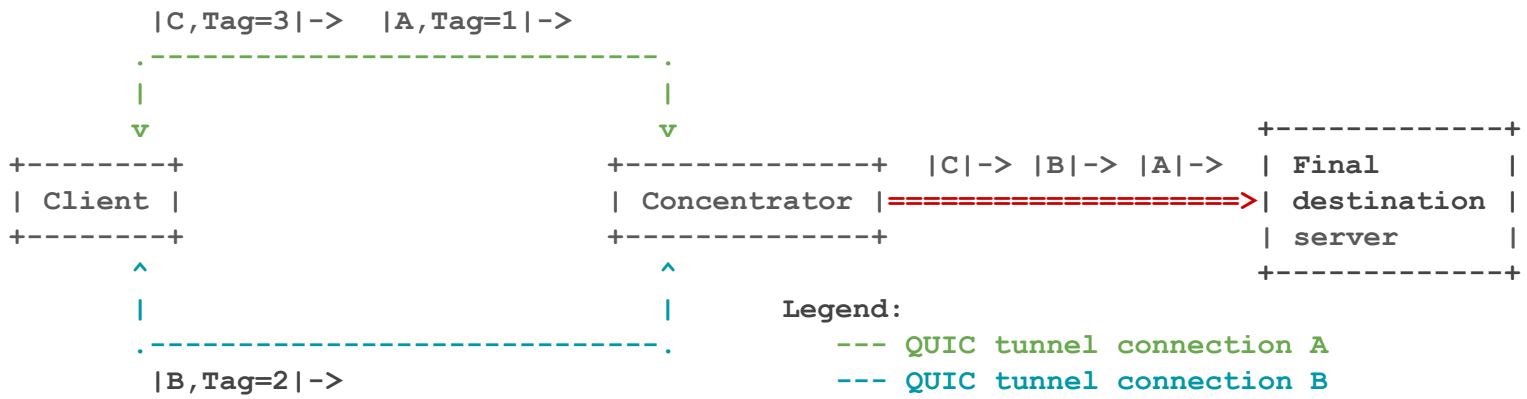

Content

- The tunnel mode
- The tunnel session mode

  draft-piraux-intarea-quic-tunnel-session
- **The stream mode**

  draft-piraux-intarea-quic-tunnel-tcp
- Conclusion

Encapsulation overhead

- Many protocols can be conveyed using these approaches.
- But it implies a significant byte overhead.
- We introduce another operating mode dedicated for conveying TCP bytestreams.

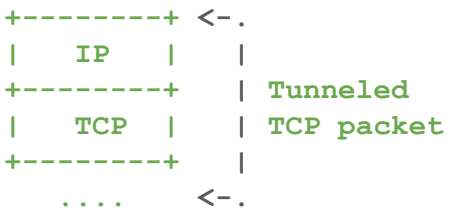

Encapsulation overhead

- Many protocols can be conveyed using these approaches.
- But it implies a significant byte overhead.
- We introduce another operating mode dedicated for conveying TCP bytestreams.

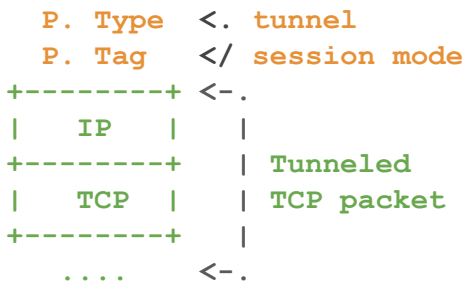

Encapsulation overhead

- Many protocols can be conveyed using these approaches.
- But it implies a significant byte overhead.
- We introduce another operating mode dedicated for conveying TCP bytestreams.

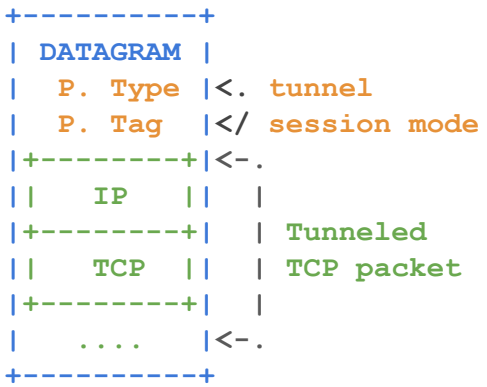

Encapsulation overhead

- Many protocols can be conveyed using these approaches.
- But it implies a significant byte overhead.
- We introduce another operating mode dedicated for conveying TCP bytestreams.

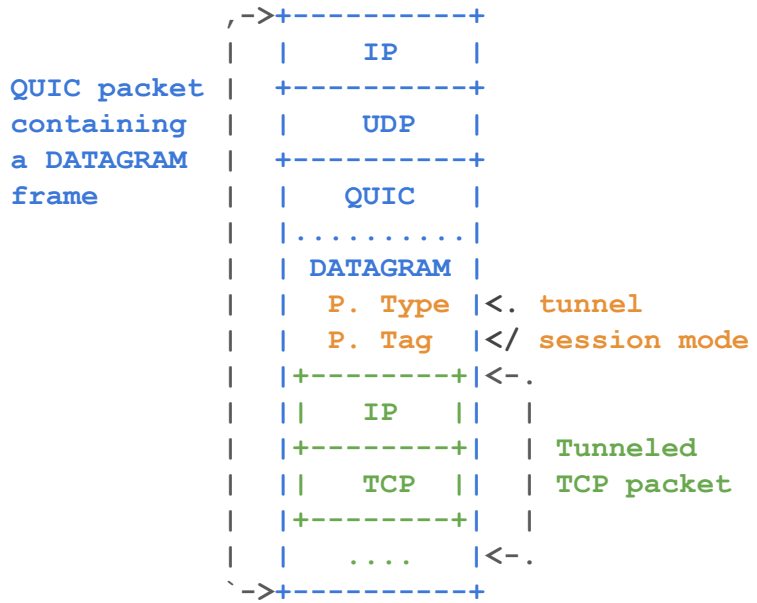

The stream mode

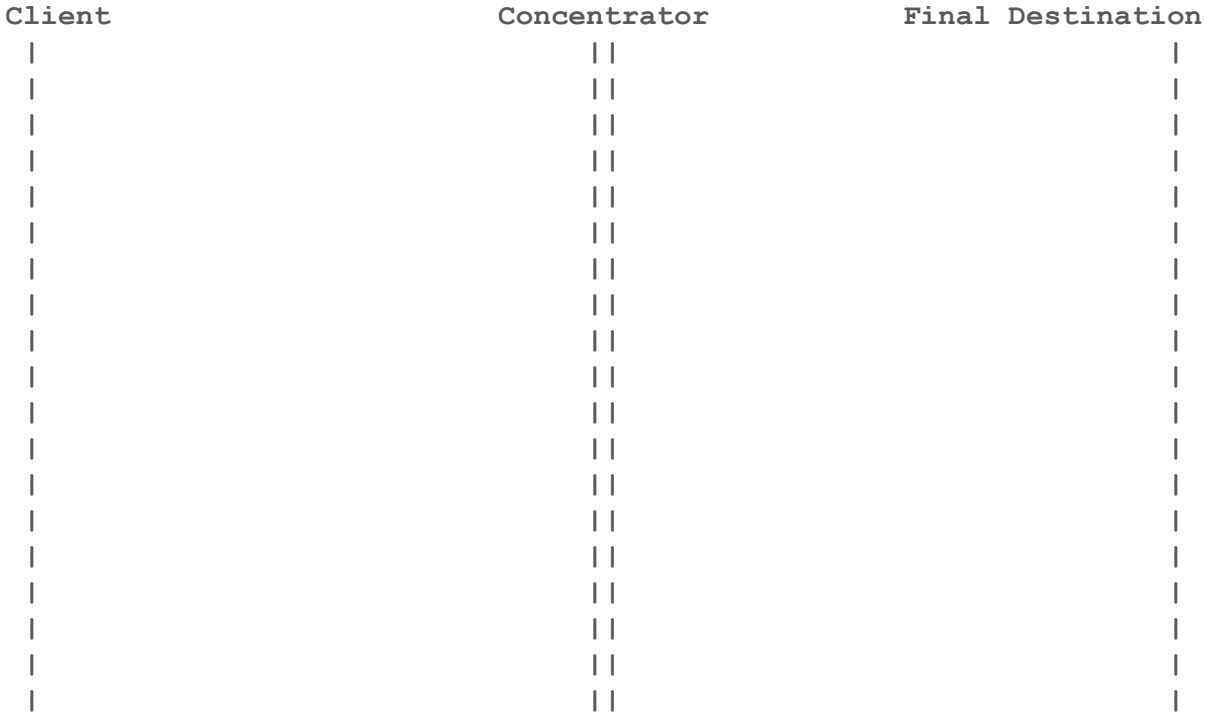

- TCP connection to QUIC stream mapping

The stream mode

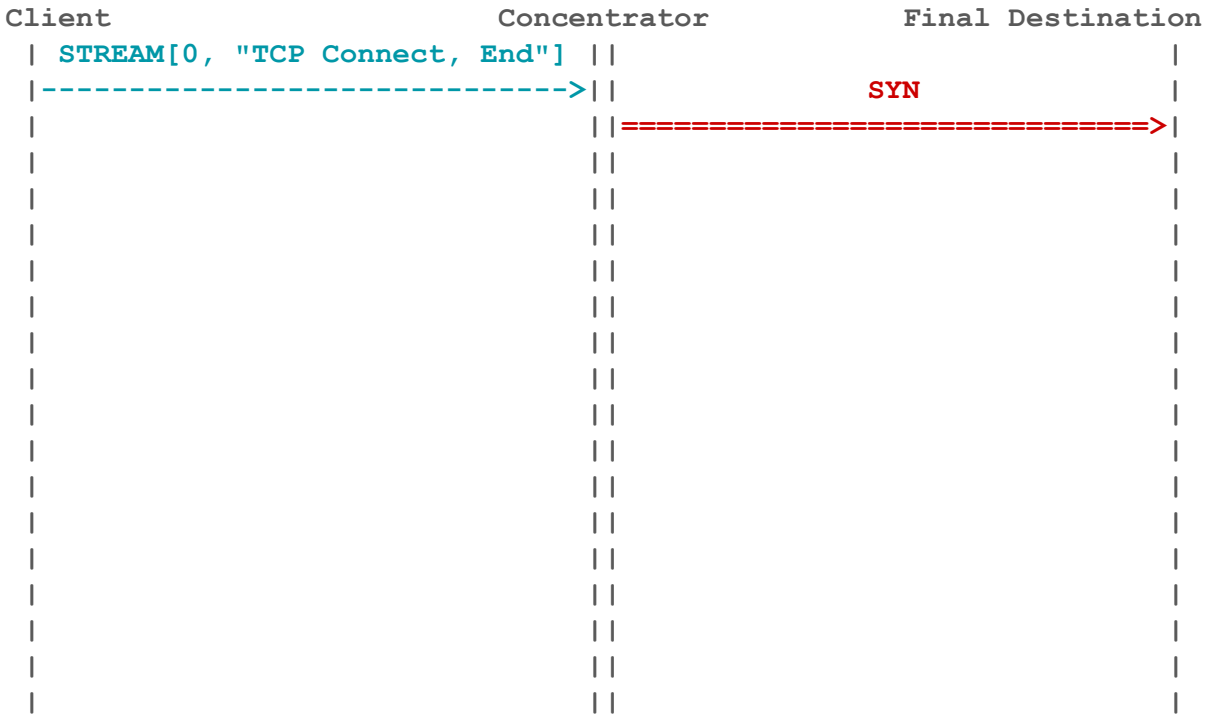

- TCP connection to QUIC stream mapping
- Client initiates a connection with a TLV

The stream mode

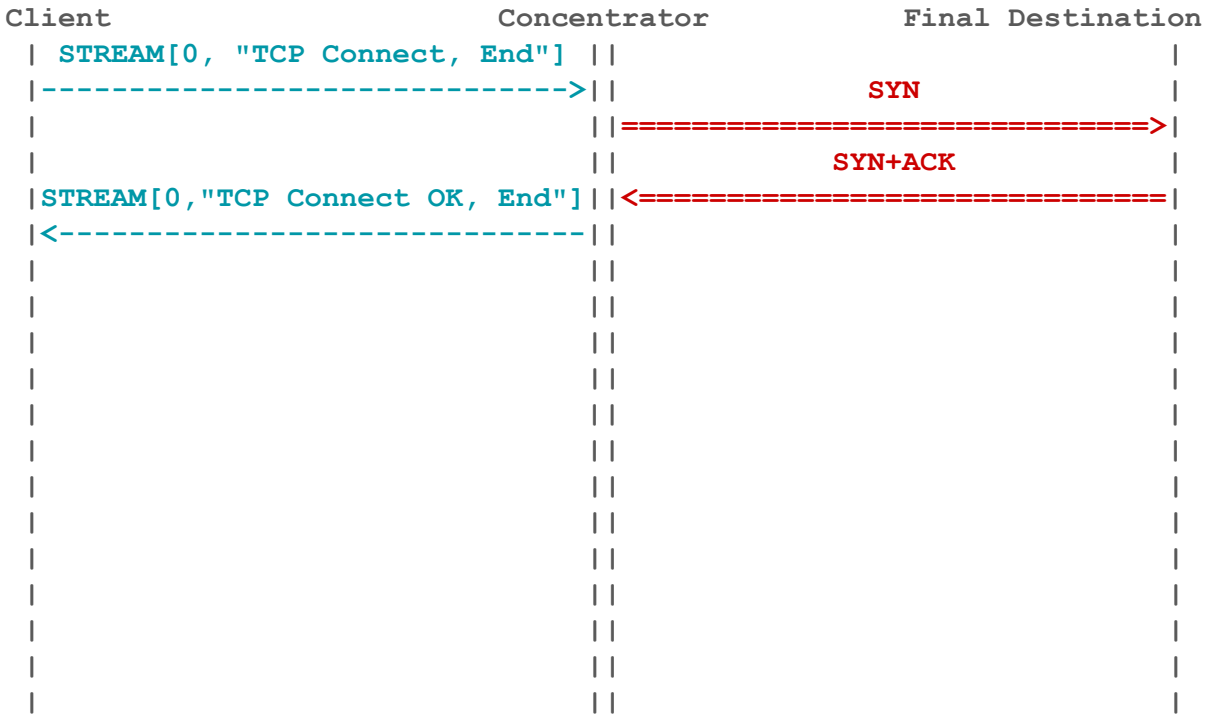

- TCP connection to QUIC stream mapping
- Client initiates a connection with a TLV

The stream mode

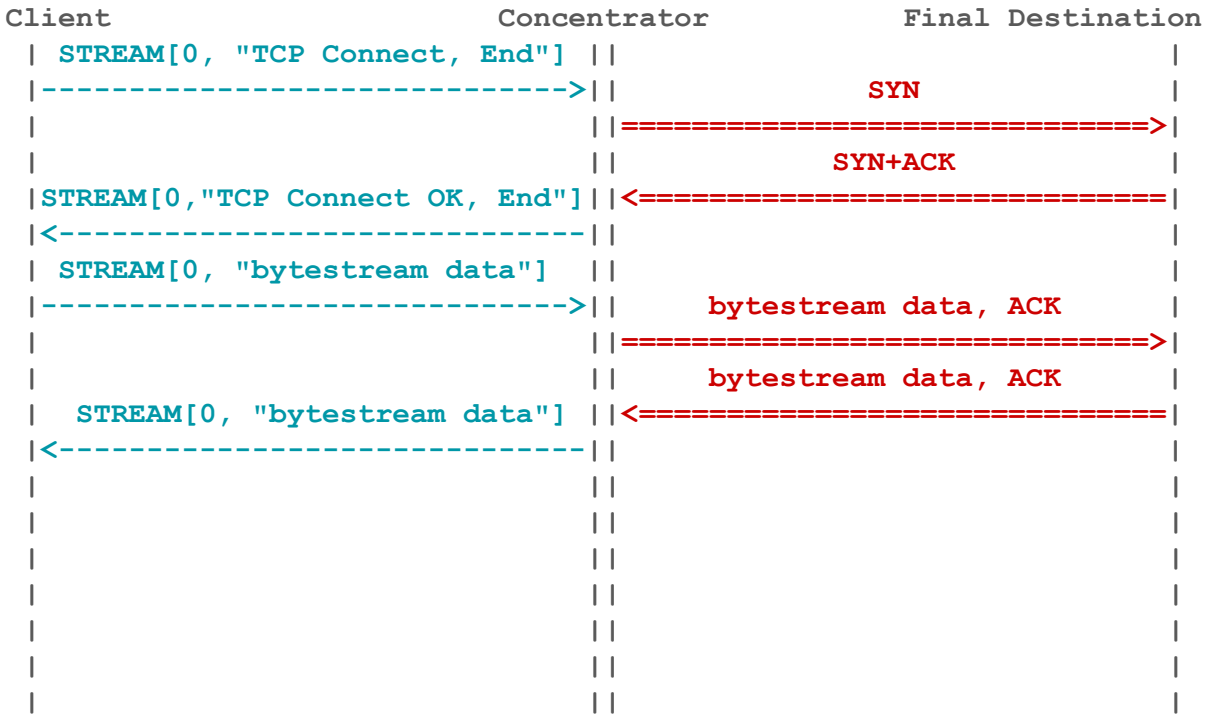

- TCP connection to QUIC stream mapping
- Client initiates a connection with a TLV
- The TCP bytestream is then copied to the QUIC stream

The stream mode

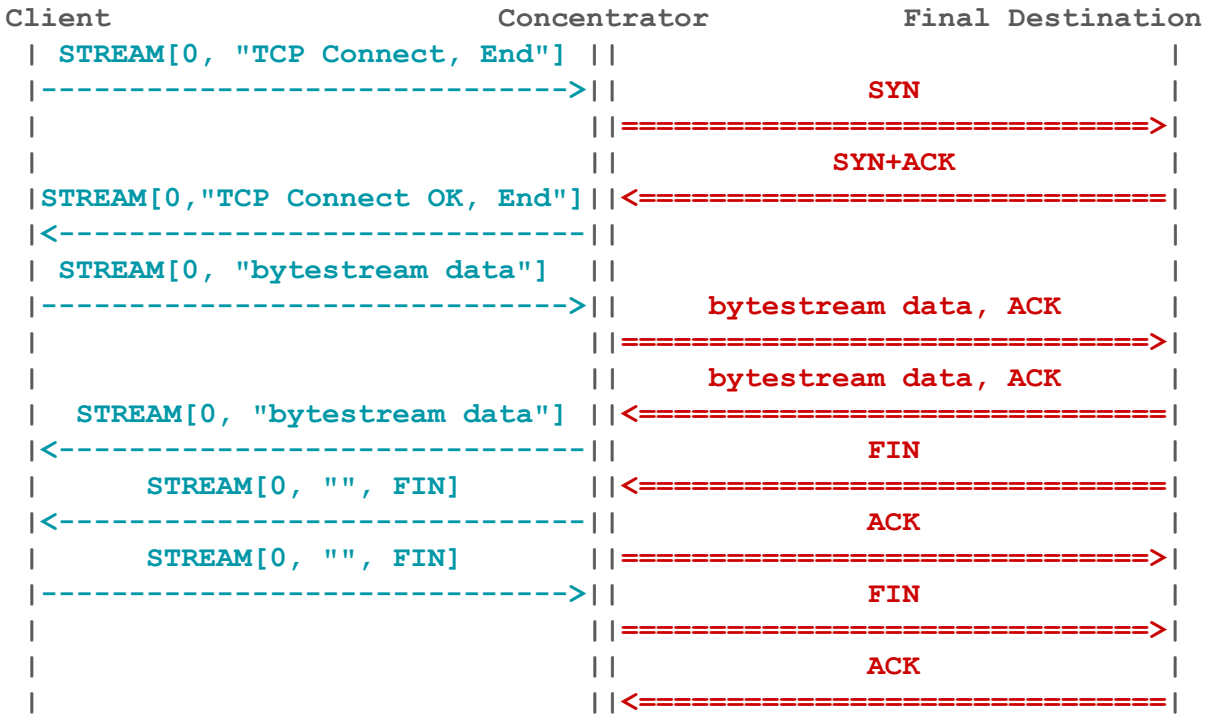

- TCP connection to QUIC stream mapping
- Client initiates a connection with a TLV
- The TCP bytestream is then copied to the QUIC stream

Conclusion

- QUIC can be used to convey many network protocols efficiently.
- There is an interest in considering multihomed clients from the start.
- We defined an application protocol to convey Internet protocols inside QUIC.
  - draft-piraux-intarea-quic-tunnel
  - draft-piraux-intarea-quic-tunnel-session
  - draft-piraux-intarea-quic-tunnel-tcp
- A partial prototype exists as part of **[PQUIC]**, see pquic.org.
- Contributions and collaborations are welcomed at github.com/mpiraux/draft-piraux-quic-tunnel.

[PQUIC] Pluginizing QUIC, *Q. De Coninck, F. Michel, M. Piraux et al.*, SIGCOMM'19

MASQUE

- Protocol defined atop HTTP/3.
- “Impacts on address migration, NAT rebinding, and future multipath mechanisms of QUIC are not anticipated”.

QUIC Tunnel

- Simple binary protocol atop QUIC.
- Considers multihomed devices from the start.

The stream mode

TCP	QUIC Stream
SYN received	Open Stream, send TLVs
FIN received	Send Stream FIN
RST received	Send STOP_SENDING
	Send RESET_STREAM
Data received	Send Stream data

QUIC Stream	TCP
Stream opened, TLVs received	Send SYN
Stream FIN received	Send FIN
STOP_SENDING received	Send RST
RESET_STREAM received	Send RST
Stream data received	Send data

- We propose a one-to-one mapping between a TCP connection and a QUIC stream.
- The Client initiates QUIC streams with a special TLV indicating the final destination.
- Then the TCP bytestream is copied to the QUIC stream data.
- A TLV for indicating a connection failure also exist.